

```

DO K - 1 TO ZTDROWS
  ADDRESS "ISPEXEC" "TBSKIP OBJLIST"
  QUEUE "-STA DB("DBNAME") SPACE("SPNAME") ACCESS(FORCE)"
  QUEUE "END"
  "DSN SYSTEM("ZZSSID")"
END
RETURN;

```

CONCLUSION

This simple function has been a significant time-saver. In the real world of DB2 where the DBA knows best when a CHECK DATA or image copy is really necessary and when it is not, the ability to override the pending statuses of DB2 objects in an expedient manner can be a real productivity boost.

Tom Sager
Systems Programmer
Louisville Gas and Electric (USA)

© Xephon 1994

New system catalog statistics in DB2 Version 3

INTRODUCTION

Most new versions and releases of DB2 cause changes to be made to the system catalog tables. Sometimes the changes are minor, sometimes they are more significant. DB2 Version 3 is an example of a product release that brought about significant changes to the system catalog.

Five new catalog tables were added for DB2 Version 3, four of which are used to house parallel statistics and one which remedies past sins. Additionally, as with most DB2 releases, there were numerous additional columns added throughout the various catalog tables.

The additional catalog information breaks down into the following basic categories:

- Partition-level statistics for query I/O parallelism
- Enhanced non-uniform distribution statistics
- Data compression statistics
- Timestamp information
- Miscellaneous.

PARTITION-LEVEL STATISTICS

DB2 Version 3 introduces partition-level statistics for many of the familiar access path related statistics, such as FREQUENCY, HIGH2KEY, LOW2KEY, NLEAF, NEVELS, NACTIVE, etc. These statistics are collected and maintained in four new tables that are similar to the catalog tables holding tablespace and indexspace aggregate statistics. Notice the new terminology: the statistics that were previously collected are now referred to as aggregate statistics. This was done because even partitioned tablespaces and indexspaces will have aggregate statistics stored in SYSTABLESPACE and SYSINDEXES. The partition-level statistics will be stored in new catalog tables as detailed below.

SYSIBM.SYSCOLSTATS

Contains one row of general partition-level statistics for each column specified to RUNSTATS.

- 1 HIGHKEY – CHAR(8)
- 2 HIGH2KEY – CHAR(8)
- 3 LOWKEY – CHAR(8)
- 4 LOW2KEY – CHAR(8)
- 5 COLCARD – INTEGER
- 6 STATSTIME – TIMESTAMP
- 7 IBMREQD – CHAR(1)
- 8 PARTITION – SMALLINT
- 9 TBOWNER – CHAR(8)
- 10 TBNAME – VARCHAR(18)
- 11 NAME – VARCHAR(18)
- 12 COLCARDATA – VARCHAR(1000).

Index

- DSNTCX01 U
TBOWNER.TBNAME.NAME
PARTITION.

Relationship

- DSND@TC
TBOWNER.TBNAME.NAME
References SYSCOLUMNS.

SYSIBM.SYSCOLDISTSTATS

Contains a maximum of ten rows of partition-level non-uniform distribution statistics for the key column of each partitioned index.

- 1 FREQUENCY – SMALLINT
- 2 STATSTIME – TIMESTAMP
- 3 IBMREQD – CHAR(1)
- 4 PARTITION – SMALLINT
- 5 TBOWNER – CHAR(8)
- 6 TBNAME – VARCHAR(18)
- 7 NAME – VARCHAR(18)
- 8 COLVALUE – VARCHAR(254).

Index

- DSNTPX01 D
TBOWNER.TBNAME.NAME.
PARTITION.

Relationship

- DSND@TP
TBOWNER.TBNAME.NAME
References SYSCOLUMNS.

SYSIBM.SYSINDEXSTATS

Contains one row of partition-level statistics for each index partition.

- 1 FIRSTKEYCARD – INTEGER
- 2 FULLKEYCARD – INTEGER

- 3 NLEAF – INTEGER
- 4 NLEVELS – SMALLINT
- 5 IOFACTOR – SMALLINT
- 6 PREFETCHFACTOR – SMALLINT
- 7 CLUSTERRATIO – SMALLINT
- 8 STATSTIME – TIMESTAMP
- 9 IBMREQD – CHAR(1)
- 10 PARTITION – SMALLINT
- 11 OWNER – CHAR(8)
- 12 NAME – VARCHAR(18)
- 13 KEYCOUNT – INTEGER.

Index

- DSNTXX01 U
OWNER.NAME.PARTITION.

Relationship

- DSNDX@TX
OWNER.NAME
References SYSINDEXES.

SYSIBM.SYSTABSTATS

Contains one row of partition level statistics for each tablespace partition.

- 1 CARD – INTEGER
- 2 NPAGES – INTEGER
- 3 PCTPAGES – SMALLINT
- 4 NACTIVE – INTEGER
- 5 PCTROWCOMP – SMALLINT
- 6 STATSTIME – TIMESTAMP
- 7 IBMREQD – CHAR(1)
- 8 DBNAME – CHAR(8)
- 9 TSNAME – CHAR(8)
- 10 PARTITION – SMALLINT
- 11 OWNER – CHAR(8)
- 12 NAME – VARCHAR(18).

Index

- DSNTTX01 U
OWNER.NAME.PARTITION.

Relationship

- DSNDT@TT
OWNER.NAME
References SYSTABLES.

After statistics have been collected for each partition, they are used to build the aggregate statistics in the tablespace and indexspace catalog tables.

DDL can be used to INSERT, UPDATE, and DELETE rows from SYSCOLDISTSTATS. Rows in SYSINDEXSTATS and SYSTABSTATS can be updated (but not inserted or deleted).

The partition-level statistics will be used to determine the number of parallel I/O streams used to access data. A preliminary number of I/O streams is ascertained at bind time. This number can be adjusted later as the query executes.

During the BIND process, DB2 reads the partition statistics from the catalog, estimates CPU cycles and I/O costs for processing the partitions, and determines the optimal number of parallel I/O streams to minimize overall elapsed time. The planned number of parallel I/O streams is also known as the DEGREE. If EXPLAIN was set to YES, the PLAN_TABLE will contain the planned degree of parallelism for the query in the ACCESS_DEGREE and JOIN_DEGREE columns. The degree can be changed during execution based upon resource availability (memory, MVS/ESA sort) and program particulars (host variable values and ambiguous cursors). Field QW0221AD in IFCID 0221 contains the actual number of parallel I/O streams used for a particular query.

NON-UNIFORM DISTRIBUTION STATISTICS (NUDS)

As of DB2 2.3, the SYSIBM.SYSFIELDS table was not in first normal form because it contained non-atomic data. Some rows contained information about field procedures, whereas other rows contained Non-Uniform Distribution Statistics (NUDS). The contents of each

row was determined by the data values contained in specific columns of the table. This was a severe design flaw that should never have been allowed into DB2 in the first place. This anomaly did not occur until DB2 2.2, when non-uniform distribution statistics were first added to the DB2 catalog.

As of DB2 Version 3, aggregate non-uniform distribution statistics are collected in a new table called SYSCOLDIST, instead of SYSFIELDS.

SYSIBM.SYSCOLDIST

Contains up to ten non-uniform distribution statistics rows for each high order index column.

- 1 FREQUENCY – SMALLINT
- 2 STATTIME – TIMESTAMP
- 3 IBMREQD – CHAR(1)
- 4 TBOWNER – CHAR(8)
- 5 TBNAME – VARCHAR(18)
- 6 NAME – VARCHAR(18)
- 7 COLVALUE – VARCHAR(254).

Index

- DSNTNX01 D
TBOWNER.TBNAME.NAME.

Relationship

- DSNDC@TN
TBOWNER.TBNAME.NAME
References SYSCOLUMNS.

Old NUDS currently stored in SYSFIELDS will remain there until RUNSTATS is executed again, or until the MODIFY STATISTICS utility is executed to remove the non-uniform distribution rows.

A row is inserted into SYSCOLDIST for each of up to ten most frequently occurring first key column index values. Similarly, a row is inserted in the SYSCOLDISTSTATS table for up to ten of the most frequently occurring first key column index values per partition, along with rows for the high and low first key index values.

DDL can be used to INSERT, UPDATE, and DELETE rows from SYSCOLDIST.

DATA COMPRESSION STATISTICS

DB2 Version 3 comes with built-in hardware assisted compression. To administer and manage compressed data, additional statistics have been added to the DB2 catalog. Data compression information columns were added to SYSTABLEPART and SYSTABLES.

The PAGESAVE column has been added to SYSTABLEPART. It indicates the percentage of pages in a partition or tablespace that was saved by using data compression. The additional bytes required for the compression dictionary used by the Ziv-Lempel compression algorithm are factored into the calculation for this number. The value can be a negative number indicating that compression actually requires additional storage space. A value of 0 indicates that either this statistic was not updated (ie RUNSTATS was not run) or that data compression resulted in no page savings.

Also added to the SYSTABLEPART table is the COMPRESS column. It simply indicates whether compression is defined for the table identified by the row.

Another final new data compression-related column, PCTROWCOMP, was added to SYSTABLES. It shows the percentage of active rows, for either a partition or table, that were affected by either compression or an EDITPROC that shortened the length of the row. A value of -1 indicates that either this statistic has not been updated or the row in the table is for a view or alias. PCTROWCOMP is user-updateable.

TIMESTAMP STATISTICS

A new column, STATSTIME, has been included in many of the catalog tables to indicate the date and time that statistics in these tables were last updated by RUNSTATS. The same value, which is displayed in the RUNSTATS execution output, is used for every STATSTIME column affected by a single run of the utility.

The catalog tables that contain this new column are:

- SYSCOLUMNS
- SYSCOLDIST
- SYSCOLDISTSTATS
- SYSCOLSTATS
- SYSINDEXES
- SYSINDEXPART
- SYSINDEXSTATS
- SYSTABLEPART
- SYSTABLES
- SYSTABLESPACE
- SYSTABSTATS.

The STATTIME column can be updated with DDL in all of the catalog tables that contain statistics used by the DB2 optimizer.

MISCELLANEOUS STATISTICS

Several other miscellaneous changes were made to the DB2 catalog as outlined below:

- The SPACE column was added to SYSTABLEPART to house partition-level DASD usage information. This information is collected by the STOSPACE utility.
- The STYPE and PIT_RBA columns were added to SYSCOPY to store enhanced information about image copy datasets. STYPE houses utility termination information and PIT_RBA houses point-in-time recovery information.
- Two new columns were added to SYSINDEXES, namely IOFACTOR and PREFETCHFACTOR. However, they are not currently used by DB2 Version 3, but are reserved for future use.
- The INDEXTYPE column was added to SYSINDEXPART, but this column will not be used until DB2 Version 4, when the new index manager feature provides a new type of index. For now, DB2 Version 3 does not utilize this column.
- Both SYSPLAN and SYSPACKAGE were modified to include the DEGREE column. It specified the number of parallel I/O streams that were determined at bind time for the plan or package.

- **SYSPPLAN** also added two additional columns: **SQLRULES** and **DISCONNECT**. **SQLRULES** indicates the type of connections to be used, ie standard (STD) or DB2. The **DISCONNECT** option specifies which remote connections should end after the completion of a logical unit of work. Three options exist, **AUTOMATIC**, **CONDITIONAL**, and **EXPLICIT**.
- Finally, the **LOCKMAX** column was added to **SYSTABLESPACE**. As of Version 3, it is not being used.

BENEFITS OF THE STATISTICAL CHANGES

There are many benefits that can be derived from DB2 storing the additional catalog information. Primarily they will be used to influence performance, management, and database administration.

The partition-level statistics in the **SYSCOLDIST**, **SYSCOLDISTSTATS**, and **SYSTABSTATS** catalog tables are used to determine the degree of parallelism. Additionally, these statistics, along with those in **SYSINDEXSTATS**, are used to build aggregate statistics in **SYSCOLDIST**, **SYSCOLUMNS**, **SYSINDEXES**, and **SYSTABLES**, which are also used by the optimizer. The new **NUDS** are also used by the optimizer to determine filter factors.

PCTROWCOMP, in **SYSTABLES**, is used by the optimizer when estimating the amount of CPU resources needed for each possible access path. **PCTROWCOMP** in **SYSTABSTATS** is used to build an aggregate **PCTROWCOMP** value in **SYSTABLES** when a tablespace is partitioned.

From a management and administration perspective, consider the following:

- Data compression can be monitored using **PAGESAVE**. If this value decreases over time, the compression dictionary may need to be rebuilt or the nature of the data may have changed rendering compression less effective.
- The **STATSTIME** information stored in the various system catalog tables may be used to determine the last time the **RUNSTATS** was used to update the statistics.

- Partition statistics in SYSINDEXSTATS can be used to determine when a specific index partition should be reorganized.
- Since some of the new statistical columns and tables are updateable, statistical values can more easily be migrated between subsystems. For instance, statistics could be updated in a test subsystem to mirror those in a production subsystem, thus enabling the programmers to predict access path selection in the production subsystem while using the test subsystem.

IMPLICATIONS OF ADDITIONAL STATISTICS

The additional catalog statistics also impact other system components. For example, the RUNSTATS utility can be executed on a partition by partition basis using the new PART parameter. The statistics can be collected at the data or index partition level. This may have ramifications for the manner in which RUNSTATS is to be issued at your shop. Instead of running RUNSTATS monthly against an entire tablespace, it may be wise to run it a partition at a time on a weekly basis.

The manner in which NUDS are maintained will change. Prior to DB2 Version 3 the MODIFY STATISTICS utility was required to remove NUDS from the catalog. After migrating to DB2 Version 3, SYSCOLDIST can be managed using DDL instead. In fact, MODIFY STATISTICS will not work on SYSCOLDIST and SYSCOLDISTSTATS. The Version 3 migration job moves non-uniform distribution statistics from SYSFIELDS to SYSCOLDIST. Once there, these statistics are maintained through SQL statements, not, as previously done, by the MODIFY STATISTICS utility.

The most significant change, though, is the additional flexibility afforded to the DBA to manipulate system catalog statistics using DDL. This enables:

- Easier population of production statistics to a test environment.
- Easier performance tweaking by updating catalog statistics (even though this method is not generally recommended).

MIGRATION CONSIDERATIONS

Because new statistics are available, access paths could change for packages and plans if they are rebound under Version 3. If it becomes necessary to fall back to DB2 2.3, statistics in SYSFIELDS must be recreated by running RUNSTATS. These statistics are not copied from SYSCOLDIST to SYSFIELDS as part of the fallback procedure. If data is compressed, it must be decompressed during the execution of RUNSTATS.

SYNOPSIS

DB2 Version 3 provides a wealth of new statistics that need to be properly understood in order to successfully manage a DB2 Version 3 implementation.

Craig S Mullins
Senior Technical Advisor
Platinum Technology (USA)

© Craig S Mullins 1994

DB2 Version 4

BACKGROUND

IBM unveiled the contents of Version 4 of DB2 (or DB2 for MVS/ESA as it is now styled) on 27 September 1994 at its DB2 Technical Conference in Orlando, Florida. Particular emphasis has been given to increased functionality and exploitation of the latest parallel technology. New facilities can be categorized in terms of improvements and enhancements to availability, client/server support, usability, and performance.

AVAILABILITY

With more and more I/S managers being asked to deliver full 24 x 7 (x 52?) availability, DB2 developers have had to look for ways of allowing concurrent access to data at all times. IBM's implementation